

POLYNOMIAL CONSTRAINT SATISFACTION PROBLEMS, GRAPH BISECTION, AND THE ISING PARTITION FUNCTION

ALEXANDER D. SCOTT* AND GREGORY B. SORKIN

ABSTRACT. We introduce a problem class we call Polynomial Constraint Satisfaction Problems, or PCSP. Where the usual CSPs from computer science and optimization have real-valued score functions, and partition functions from physics have monomials, PCSP has scores that are arbitrary multivariate formal polynomials, or indeed take values in an arbitrary ring.

Although PCSP is much more general than CSP, remarkably, all (exact, exponential-time) algorithms we know of for 2-CSP (where each score depends on at most 2 variables) extend to 2-PCSP, at the expense of just a polynomial factor in running time for adding and multiplying polynomials rather than reals. Specifically, four algorithms that extend are: the reduction-based algorithm of Scott and Sorkin; the specialization of that approach to sparse random instances, where the algorithm runs in polynomial expected time; dynamic-programming algorithms based on tree decompositions; and the split-and-list matrix-multiplication algorithm of Williams.

This gives the first polynomial-space exact algorithm more efficient than exhaustive enumeration for the well-studied problems of finding a minimum bisection of a graph, and calculating the partition function of an Ising model. Furthermore, PCSP solves both optimization and counting versions of a wide range of problems, including all CSPs, and thus enables samplers including uniform sampling of optimal solutions and Gibbs sampling of all solutions.

1. INTRODUCTION

We introduce a problem class we call Polynomial 2-CSP, which generalizes more common 2-CSP constraint satisfaction problems, yet is solvable by adaptations of all the best CSP algorithms we know of, with essentially equal efficiency. This provides the first efficient polynomial-space exact algorithm for the well-studied problem of finding the maximum bisection of a graph (exponential-space algorithms were previously known), and the first exact algorithm of any sort, other than naive exhaustive enumeration, for the equally important problem of computing the partition function of an Ising model. It also solves a breadth of other optimization problems beyond the class 2-CSP, including some quite convoluted ones, in a simple, uniform framework. It allows counting and sampling as well as optimization, and again, does this in time bounds matching those for 2-CSP. Let us elaborate.

Many decision and optimization problems, such as 3-colorability, Max Cut, Max k -Cut, and Max 2-Sat, belong to the class Max 2-CSP of constraint satisfaction problems with at most two variables per constraint. A Max 2-CSP (or for short, simply CSP) instance defines a “score function” or “soft constraint” on each vertex and edge of a “constraint graph”, and its solution is a vertex coloring or “assignment” maximizing the total score.

We define a more general class we call Polynomial 2-CSP, or simply 2-PCSP or PCSP. (The PCSP definition extends obviously to longer clauses, but much less is known about algorithms even for 3-CSP.) Comparing with a CSP, the vertex and edge scores of a PCSP are (potentially multivariate) polynomials (rather than real numbers), the score of a coloring is the product of all vertex and edge scores (rather than their sum), and an instance’s “value”, or “partition function”, is the sum of scores over all colorings (rather than the maximum). In fact the scores may be “generalized polynomials”, whose powers are not restricted to positive integers but may be arbitrary reals, but we use “polynomial” for want of a better term.

* Research supported in part by EPSRC grant GR/S26323/01.

Compared with the “max–sum” CSP formulation typical in the CSP community, PCSP has a “sum–product” form commonly seen as the “partition function” in statistical physics: a sum over configurations, of a product of exponentials of “energies”. In the physics settings we are familiar with, the partition function is always a product of monomials (sometimes multivariate, for example with energy and magnetization), whereas our PCSPs permit polynomials. With general polynomials, PCSP is closed under a variety of operations enabling its (relatively) efficient solution by a number of algorithmic approaches.

PCSP includes many optimization problems not in the class CSP, including Maximum Bisection, Maximum Clique, Sparsest Cut, various judicious partitioning problems, Max Ones 2-Sat, and many others. Capitalizing on the fact that a PCSP’s score functions may be multivariate polynomials, in all these cases, we simply track several score functions at once. For example, for Max Bisection we track the size of one partition as well as the size of the cut.

The H -coloring of a graph G has received much recent attention, including the book [HN04], the application to soft constraints in [BG05], and the application to asymmetric hard constraints in [DGP06]. H -coloring can be encoded as a PCSP by giving every edge of G the same (symmetric) score function. However, PCSP can also give a different (and asymmetric) score function to each edge, and indeed this is an essential feature as it leaves the class closed under various reductions.

PCSPs give a powerful technique for counting and sampling solutions. Given a CSP, there is a natural way to obtain a corresponding PCSP. Solving the PCSP gives a generating function for the number of CSP assignments of each possible score, which allows us to count or (by successively solving subinstances) to sample uniformly from solutions of maximum score, or to sample all solutions according to the Gibbs distribution or any other score-based distribution.

What makes PCSP so interesting is that, despite being more general than CSP, all the best CSP algorithms we know of can be extended to solve PCSP, with the same running-time bounds. (Exact CSP algorithms are generally exponential-time, and solving a PCSP takes just a polynomial factor longer than solving a CSP instance of the same size, for the extra overhead of adding and multiplying polynomials rather than reals.) Specifically, four algorithms that extend are: the split-and-list matrix-multiplication algorithm of Williams [Wil04]; the reduction-based algorithm of Scott and Sorkin [SS07]; the specialization of that approach to sparse random instances, where the algorithm runs in polynomial expected time [SS06b]; and dynamic-programming algorithms based on tree decompositions [JKLS05, SS07, KMRR09].

Both the definition of PCSP and all our algorithms only require us to take sums and products of polynomial “scores”. It is thus natural to work in a still more general class that we call Ring CSP, where each score function and the “value” of an instance are given by element of an arbitrary ring R ; PCSPs are a special case where R is a polynomial ring.

2. OUTLINE

Following a small amount of notation in Section 3, Section 4 defines the classes CSP, PCSP and RCSP. Section 5 provides examples of CSPs and related PCSPs; because this includes many problems beyond CSP, such as graph bisection, the Ising partition function, and judicious partitioning, it is a focal point of the paper.

In Sections 6–9, we show how four CSP algorithms extend to PCSP/RCSP, and analyze them in terms of the number of ring operations (or sums and products of polynomials in the case of PCSPs) that they require. In Section 10 we consider the time and space complexity of the ring operations for PCSPs. We put the two together in Section 11, summarizing the time and space complexity of each algorithm, and noting the circumstances favoring one algorithm over another.

In addition to counting maximum CSP solutions, a PCSP’s partition function also enables construction of an optimal solution, and various forms of perfect random sampling from optimal solutions or all solutions; this is taken up in Section 12.

3. NOTATION

In the next section we will define the class of CSPs and our new class PCSP of Polynomial CSPs. An instance of either has a “constraint graph” $G = (V, E)$ with vertex set V and edge set E , and we reserve the symbols G , V and E for these roles. An instance of CSP or PCSP also has a domain of *values* or *colors* that may be assigned to the vertices (variables), for example $\{\text{true}, \text{false}\}$ for a satisfiability problem, or a set of colors for a graph coloring problem. In general we will denote this domain by $[k]$, interpreted as $\{0, \dots, k-1\}$ or (it makes no difference) $\{1, \dots, k\}$. At the heart of both the CSP and PCSP instance will be cost or “score” terms $s_v(i)$ ($v \in V(G)$, $i \in [k]$) and $s_{xy}(i, j)$ ($x, y \in V(G)$, $i, j \in [k]$).

For CSPs the scores s are real numbers, while for PCSPs they are polynomials which we write in the variable z if univariate, or if multivariate over z, w or z, w_1, w_2, \dots . Just as CSP scores s need not be positive integers, exponents in the PCSP score “polynomials” may be fractional or negative (or both). (The term “fractional polynomial” is used in the literature of statistical regression.)

We use \uplus to indicate disjoint union, so $V_0 \uplus V_1 = V$ means that V_0 and V_1 partition V , i.e., $V_0 \cap V_1 = \emptyset$ and $V_0 \cup V_1 = V$. The notation O^* hides polynomial factors in any parameters, so for example $O(n2^{5+19m/100})$ is contained in $O^*(2^{19m/100})$.

4. CSP, POLYNOMIAL CSP, AND RING CSP

4.1. CSP. Let us begin by defining the problem class CSP over a domain of size k . An *instance* I of CSP with constraint graph $G = (V, E)$ and domain $[k]$ has the following ingredients:

- (1) a real number s_\emptyset ;
- (2) for each vertex $v \in V$ and color $i \in [k]$, a real number $s_v(i)$;
- (3) for each edge $xy \in E$ and any colors $i, j \in [k]$, a real number $s_{xy}(i, j)$.

We shall refer to these quantities as, respectively, the *nullary score*, the *vertex scores*, and the *edge scores*. Note that we want only one score for an edge $\{i, j\}$ with given colors $\{x, y\}$ assigned its respective endpoints, so $s_{xy}(i, j)$ and $s_{yx}(j, i)$ are taken to be equivalent names for the same score (or one may simply assume that $x < y$).

Given an *assignment* (or *coloring*) $\sigma : V \rightarrow [k]$, we define the *score* of σ to be the real number

$$I(\sigma) := s_\emptyset + \sum_{v \in V} s_v(\sigma(v)) + \sum_{xy \in E} s_{xy}(\sigma(x), \sigma(y)).$$

The instance’s solution is the maximum possible score

$$(1) \quad \max_{\sigma: V \rightarrow [k]} I(\sigma) = \max_{\sigma: V \rightarrow [k]} \left(s_\emptyset + \sum_{v \in V} s_v(\sigma(v)) + \sum_{xy \in E} s_{xy}(\sigma(x), \sigma(y)) \right)$$

or the assignment σ achieving the maximum; it is straightforward to get either from the other. For obvious reasons, (1) is commonly referred to as a *max-sum* formulation.

In this paper we take a generating-function approach, among other things enabling us to count the number of solutions satisfying various properties. Given an instance I of CSP, the corresponding generating function is the polynomial

$$(2) \quad \sum_{\sigma: V \rightarrow [k]} z^{I(\sigma)}.$$

More generally, if we want to keep track of several quantities simultaneously, say $I(\sigma), J(\sigma), \dots$, we can consider a multivariate generating function $\sum_{\sigma} z^{I(\sigma)} w^{J(\sigma)} \dots$. Calculating the generating function in the obvious way, by running through all $k^{|V|}$ assignments, is clearly very slow, so it is desirable to have more efficient algorithms.

In working with generating functions, we borrow some notions from statistical physics. We think of the score $I(\sigma)$ as a “Hamiltonian” measuring the “energy” of a configuration σ . Thus edge scores

correspond to “pair interactions” between adjacent sites, while vertex scores measure the effect of a “magnetic field”. (From this perspective, the nullary score is just a constant that disappears after normalization.) The generating function is then the “partition function” for this model. (The partition function is often written in the form $\sum_{\sigma} \exp(-\beta I(\sigma))$, where β is known as the *inverse temperature*, but substituting z for $e^{-\beta}$ yields the partition function in polynomial form.)

A crucial element of our approach is that the score $I(\sigma)$ can be broken up as a sum of local interactions, and thus an expression such as $z^{I(\sigma)}$ can be expressed as a product of monomials corresponding to local interactions:

$$z^{I(\sigma)} = z^{s_{\emptyset}} \cdot \prod_{v \in V} z^{s_v(\sigma(v))} \cdot \prod_{xy \in E} z^{s_{xy}(\sigma(x), \sigma(y))}.$$

In order to provide a framework for this approach, we introduce a generalized version of constraint satisfaction where the scores are polynomials (in some set of variables) instead of real numbers, and the score of an assignment is taken as a product rather than a sum.

4.2. Polynomial CSP. An instance I of Polynomial CSP, with constraint graph G and domain $[k]$, has the following ingredients:

- (1) a set of formal variables, and various polynomials over these variables;
- (2) a polynomial p_{\emptyset} ;
- (3) for each vertex $v \in V$ and color $i \in [k]$, a polynomial $p_v(i)$;
- (4) for each edge $xy \in E$ and pair of colors $i, j \in [k]$, a polynomial $p_{xy}(i, j)$.

We refer to these three types of polynomial as, respectively, the *nullary polynomial*, the *vertex polynomials*, and the *edge polynomials*. (Reiterating from Section 3, we use “polynomial” in a general sense, allowing negative and fractional powers.) We want only one polynomial for a given edge with given colors on its endpoints, so again we either take $p_{xy}(i, j)$ and $p_{yx}(j, i)$ to be equivalent or simply assume that $x < y$.

Given an assignment $\sigma : V \rightarrow [k]$, we define the *score* of σ to be the polynomial

$$(3) \quad I(\sigma) := p_{\emptyset} \cdot \prod_{v \in V} p_v(\sigma(v)) \cdot \prod_{xy \in E} p_{xy}(\sigma(x), \sigma(y)).$$

We then define the *partition function* Z_I of I by

$$(4) \quad Z_I = \sum_{\sigma : V \rightarrow [k]} I(\sigma) = \sum_{\sigma : V \rightarrow [k]} p_{\emptyset} \cdot \prod_{v \in V} p_v(\sigma(v)) \cdot \prod_{xy \in E} p_{xy}(\sigma(x), \sigma(y)).$$

The partition function is the “solution” of a PCSP instance. In contrast with the max–sum formulation of (1), (4) is referred to as a *sum–product* formulation. The sense of the PCSP formulation will become clearer with examples in the next section.

4.3. Ring CSP. The definition of $I(\sigma)$ in equation (3), and for the most part our analysis in later sections, requires only addition and multiplication of scores. It therefore makes sense to work in a still more general context where the score functions are elements of an arbitrary ring. We define an instance I of RCSP over a ring R , with constraint graph G and domain $[k]$, to have the following ingredients:

- (1) a ring element r_{\emptyset} ;
- (2) for each vertex $v \in V$ and color $i \in [k]$, a ring element $r_v(i)$;
- (3) for each edge $xy \in E$ and pair of colors $i, j \in [k]$, a ring element $r_{xy}(i, j)$.

The *score* of an assignment $\sigma : V \rightarrow [k]$ is the ring element

$$I(\sigma) := r_{\emptyset} \cdot \prod_{v \in V} r_v(\sigma(v)) \cdot \prod_{xy \in E} r_{xy}(\sigma(x), \sigma(y)),$$

and the *partition function* Z_I of I is

$$Z_I = \sum_{\sigma: V \rightarrow [k]} I(\sigma).$$

Thus PCSPs are the special case of RCSPs where the ring R is a polynomial ring over the reals.

5. EXAMPLES

In this section we show how some standard problems can be written as PCSPs. First we show that for every CSP there is a naturally corresponding PCSP; then we illustrate how various problems, including some that are not CSPs, can be expressed as PCSPs. We will turn to algorithms in later sections, but for now bear in mind that a PCSP over a constraint graph G can be solved about as efficiently as a CSP over the same graph.

5.1. Generating function of a simple CSP.

Definition 1 (Generating function of an instance). *Given an instance I of CSP on a constraint graph $G = (V, E)$ and variable domain $[k]$, we can define a corresponding instance I^* of PCSP with the same graph and variable domain, and polynomials*

$$\begin{aligned} p_\emptyset &= z^{s_\emptyset} \\ (\forall v \in V, \forall i \in [k]) \quad &p_v(i) = z^{s_v(i)} \\ (\forall xy \in E, \forall i, j \in [k]) \quad &p_{xy}(i, j) = z^{s_{xy}(i, j)}. \end{aligned}$$

The connection between I and I^* is given by the following simple observation.

Lemma 2. *Let I be an instance of CSP, and let I^* be the corresponding PCSP instance. Then the partition function Z_{I^*} is the generating function (2) for the instance I .*

Proof. For any assignment σ , we have

$$\begin{aligned} I^*(\sigma) &= z^{s_\emptyset} \cdot \prod_{v \in V} z^{s_v(i)} \cdot \prod_{xy \in E} z^{s_{xy}(i, j)} \\ &= z^{s_\emptyset + \sum_{v \in V} s_v(i) + \sum_{xy \in E} s_{xy}(i, j)} \\ &= z^{I(\sigma)}. \end{aligned}$$

It therefore follows that the partition function

$$Z_{I^*} = \sum_{\sigma} I^*(\sigma) = \sum_{\sigma} z^{I(\sigma)}$$

is the generating function for the original constraint satisfaction problem. \square

Similar results are easily seen to hold for generating functions in more than one variable.

5.2. Max Cut and Max Dicut. Max Cut provides a simple illustration of Definition 1 and Lemma 2. Let us first write Max Cut as a CSP; we will then construct the corresponding PCSP.

Example 3 (Max Cut CSP). *Given a graph $G = (V, E)$, set $k = 2$ and define a CSP instance I by*

$$\begin{aligned} s_\emptyset &= 0 \\ (\forall v \in V) \quad &s_v(0) = s_v(1) = 0 \\ (\forall xy \in E) \quad &s_{xy}(0, 1) = s_{xy}(1, 0) = 1, \quad s_{xy}(0, 0) = s_{xy}(1, 1) = 0. \end{aligned}$$

With $\sigma^{-1}(i) = \{v : \sigma(v) = i\}$, note that $(V_0, V_1) = (\sigma^{-1}(0), \sigma^{-1}(1))$ is a partition of V , and

$$(5) \quad I(\sigma) = \sum_{\substack{xy \in E: \\ \sigma(x)=0, \sigma(y)=1}} 1 = e(V_0, V_1)$$

is the size of the cut induced by σ . The corresponding PCSP instance is obtained as in Definition 1.

Example 4 (Max Cut PCSP). *Given a graph $G = (V, E)$, we set $k = 2$ and define a PCSP instance I by*

$$\begin{aligned} p_\emptyset &= 1 \\ (\forall v \in V) \quad p_v(0) &= p_v(1) = 1 \\ (\forall xy \in E) \quad p_{xy}(0, 1) &= p_{xy}(1, 0) = z \quad , \quad p_{xy}(0, 0) = p_{xy}(1, 1) = 1. \end{aligned}$$

(In all such cases, a 1 on the right hand side may be thought of as z^0 .)

By Lemma 2, the partition function Z_I is therefore the generating function for cuts:

$$(6) \quad Z_I = \sum 2c_i z^i,$$

where c_i is the number of cuts of size i , and the factor 2 appears because each cut (V_0, V_1) also appears as (V_1, V_0) . The size of a maximum cut is the degree of Z_I , and the number of maximum cuts is half the leading coefficient.

Note that the partition function (6) is the partition function of the Ising model with no external field (see below for a definition). Thus we have recovered the familiar fact that, up to a change of variables, the partition function of the Ising model is the generating function for cuts.

We can also encode weighted instances of Max Cut with edge weights $w : E \rightarrow \mathbb{R}$ by modifying the third line of the definition above to

$$p_{xy}(0, 1) = p_{xy}(1, 0) = z^{w(xy)} \quad \forall xy \in E.$$

Similarly, we can encode weighted Max Dicut (maximum directed cut) by setting

$$\begin{aligned} p_\emptyset &= 1 \\ (\forall v \in V) \quad p_v(0) &= p_v(1) = 1 \\ (\forall xy \in E) \quad p_{xy}(0, 1) &= z^{w(xy)}, \quad p_{xy}(1, 0) = z^{w(yx)}, \quad p_{xy}(0, 0) = p_{xy}(1, 1) = 1, \end{aligned}$$

where $w(xy)$ denotes the weight of the *directed* edge xy , and we define $w(xy) = 0$ if there is no such edge. Max k -Cut is encoded the same way, only with k -valued variables in place of binary ones.

5.3. The Ising model and Max Bisection. A slight generalization of the previous example allows us to handle the Ising model.

Example 5 (Ising CSP model). *The Ising model with edge weights J and external field h on a graph G is defined in terms of its Hamiltonian H . For an assignment $\sigma : V \rightarrow \{0, 1\}$, we define*

$$H(\sigma) = J \sum_{xy \in E} \delta(\sigma(x), \sigma(y)) + h \sum_{v \in V} \sigma(v).$$

Here $\delta(a, b)$ is the delta function, returning 0 if $a = b$, and 1 otherwise, J is the *interaction strength*, and h is the *external magnetic field*. In analogy with (5), taking $V_i = \sigma^{-1}(i)$, we may rewrite H as

$$H(\sigma) = J e(V_0, V_1) + h |V_1|.$$

Note that H is an instance of (plain) CSP.

The *partition function* of the Ising model at *inverse temperature* β is

$$(7) \quad Z_{\text{Ising}} = \sum_{\sigma} e^{-\beta H(\sigma)} = \sum_{V_0 \uplus V_1 = V} w^{|V_1|} z^{e(V_0, V_1)},$$

where we have written $w = e^{-\beta h}$ and $z = e^{-\beta J}$, and the last sum is taken over ordered pairs (V_0, V_1) that partition V . With this change of variables, the Ising partition function is easily expressed as partition function of a PCSP, this time over two variables.

Example 6 (Ising PCSP). *Define an Ising PCSP instance I by*

$$\begin{aligned} p_{\emptyset} &= 1 \\ (\forall v \in V) \quad p_v(0) &= 1 \\ (\forall xy \in E) \quad p_v(1) &= w, \quad p_{xy}(0, 1) = p_{xy}(1, 0) = z, \quad p_{xy}(0, 0) = p_{xy}(1, 1) = 1. \end{aligned}$$

With our usual notation, $I(\sigma) = w^{|V_1|} z^{e(V_0, V_1)}$, and so the partition function for this PCSP is equal to (7).

More generally, the Potts model can be written in a similar way. Less generally, the zero magnetization ($h = 0$) anti-ferromagnetic ($J = -1$) special case of this Ising model has maximum cuts as its “ground states” (lowest-energy states): taking $\beta \rightarrow \infty$ means that these states dominate the sum comprising Z_{Ising} . To view it another way, the maximum degree in z of $Z(w, z)$ is the size of a *maximum cut* in G .

The Ising PCSP can also be used to handle *Max Bisection*. This is important, because most CSP algorithms (other than dynamic programming) can solve Max Cut, but cannot be applied to Max Bisection because there is no way to force them to generate a balanced cut. The Ising PCSP does this by tracking two variables at once: the bisections of G correspond to terms in Z_{Ising} that have degree $\lfloor n/2 \rfloor$ in w . (Each bisection is counted once if n is odd, twice if n is even.) Extracting these terms gives the generating function for bisections of G .

The same partition function also yields a *sparsest cut* of the graph: sparsest cuts correspond to terms with the largest ratio of the power of z (number of cut edges) to the power of w (number of vertices in one partition).

5.4. Max Independent Set and Max Clique. Maximum Independent Set (MIS) is easily expressed as a CSP:

Example 7 (MIS CSP).

$$\begin{aligned} s_{\emptyset} &= 0 \\ (\forall v \in V) \quad s_v(0) &= 0, \quad s_v(1) = 1 \\ (\forall xy \in E) \quad s_{xy}(0, 0) &= s_{xy}(0, 1) = s_{xy}(1, 0) = 0, \quad s_{xy}(1, 1) = -2. \end{aligned}$$

Maximum clique¹ cannot be modeled in the same way, because the clique constraint is enforced by non-edges, which are not an element of the model. Of course a maximum clique in G corresponds to a maximum independent set in its complement graph \overline{G} , but for our purposes this can be very different, as we typically parametrize running time in terms of the number of edges, and a Max Clique instance on a sparse graph G with $|E|$ edges becomes an MIS instance on the dense graph \overline{G} with $\binom{n}{2} - |E|$ edges. We discuss this further in Section 13. However, as with Max Bisection, it is possible to model a sparse Max Clique instance as a PCSP (with the input graph G as its constraint graph) by introducing a second variable.

¹We use clique in the sense of a complete subgraph, not the stronger definition as a *maximal* complete subgraph.

Example 8 (Max Clique PCSP).

$$\begin{aligned} s_\emptyset &= 1 \\ (\forall v \in V) \quad s_v(0) &= 1, \quad s_v(1) = w \\ (\forall xy \in E) \quad s_{xy}(0,0) &= s_{xy}(0,1) = s_{xy}(1,0) = 1, \quad s_{xy}(1,1) = z. \end{aligned}$$

An assignment σ has score $I(\sigma) = w^{|V_1|} z^{e(G|_{V_1})}$, the power of w counting the number of vertices in the chosen set $V_1 = \sigma^{-1}(1)$, and the power of z counting the number of edges induced by that set. In the partition function, k -cliques correspond to terms of the form $w^k z^{\binom{k}{2}}$. Of course, independent sets of cardinality k correspond to terms with $w^k z^0$: the PCSP simultaneously counts maximum cliques and maximum independent sets (among other things).

The same approach also counts maximum-weight cliques and/or independent sets, in a graph with vertex and/or edge weights, if we introduce a third formal variable w' and set the scores as before except for $s_v(1) = ww'^{\text{weight}(v)}$ and $s_{xy}(1,1) = zw'^{\text{weight}(x,y)}$. Cliques correspond to terms $w^k z^{\binom{k}{2}} w'^{\text{weight}}$, and the maximum-weight clique the term with the largest power of w' .

5.5. Judicious partitions. Similar techniques may be applied to various *judicious partitioning* problems [BS99, Sco05], such as finding a cut of a graph which minimizes $\max\{e(V_0), e(V_1)\}$.

Example 9 (Judicious bipartition PCSP).

$$\begin{aligned} s_\emptyset &= 1 \\ (\forall v \in V) \quad s_v(0) &= s_v(1) = 1 \\ (\forall xy \in E) \quad s_{xy}(0,1) &= s_{xy}(1,0) = 1, \quad s_{xy}(0,0) = z_1, \quad s_{xy}(1,1) = z_2. \end{aligned}$$

A simple calculation shows that $I(\sigma) = z_0^{e(V_0)} z_1^{e(V_1)}$, where $V_i = \sigma^{-1}(i)$. Thus the problem of minimizing $\max\{e(V_0), e(V_1)\}$ (or, for instance, the problem of maximizing $\min\{e(V_0), e(V_1)\}$) can be solved by examining the partition function.

We might further ask for a *bisection* that minimizes $\max\{e(V_0), e(V_1)\}$, and this is easily achieved by introducing a third variable: we set $s_v(0) = w$ and examine only terms of the partition function that have degree $\lfloor |V|/2 \rfloor$ in w .

5.6. Simultaneous assignments. We can think of the foregoing example, of finding a balanced bisection minimizing $\max\{e(V_0), e(V_1)\}$, as a case of having more than one CSP on a single set of variables. In other examples, we might have two instances of Max Sat that we wish to treat as a bi-criterion optimization problem; or a Max Sat instance to optimize subject to satisfaction of a Sat instance; or we might wish to find a partition of a vertex set yielding a large cut for two different graphs on the same vertices. All such problems are now straightforward: use a variable z to encode one problem and a variable w to encode the second. (Thus, the initial edge scores are of the form $z^{\text{firstscore}} w^{\text{secondscore}}$.)

6. SPLIT-AND-LIST, MATRIX-MULTIPLICATION ALGORITHM

The algorithmic approach introduced by Williams [Wil04], which he calls “split and list”, is characterized by its use of fast matrix multiplication. For binary 2-CSPs it runs in time $O^*(2^{\omega n/3})$, where ω is the “matrix multiplication exponent” such that two $n \times n$ matrices can be multiplied together in time $O(n^\omega)$. It also requires exponential space, $O(2^{2n/3})$. The best known bound on the matrix multiplication constant is $\omega < 2.37\dots$, from the celebrated algorithm of Coppersmith and Winograd [CW90], which works over arbitrary rings.

Koivisto [Koi06b] extended Williams’ algorithm from CSP (with a max–sum formulation) to counting-CSP (with sum–product formulation), and Koivisto’s version transfers immediately to the PCSP and RCSP context. While the sum–product formulation is common in statistical physics,

Koivisto's use of it appears to be one of the few besides the present work in the field of exact algorithms. For Williams' algorithm, the sum-product formulation is not only more general but also simpler.

6.1. Algorithm. Koivisto's counting extension of Williams' algorithm is very much in the spirit of the original, and itself works in the PCSP and RCSP settings. Arbitrarily partition the variables V into 3 equal-sized sets, A , B and C . Explicitly list every possible assignment $\alpha : A \mapsto [k]$ of variables in A and similarly every assignment β for B and γ for C . Explicitly construct M_{AA} , a $k^{|A|} \times k^{|A|}$ diagonal matrix with elements

$$M_{AA}(\alpha, \alpha) := \prod_{v \in A} p_v(\alpha(v)) \cdot \prod_{xy \in E \cap \{A \times A\}} p_{xy}(\alpha(x), \alpha(y)),$$

and likewise M_{BB} and M_{CC} . Also explicitly construct M_{AB} , a $k^{|A|} \times k^{|B|}$ matrix with elements

$$M_{AB}(\alpha, \beta) := \prod_{xy \in E \cap \{A \times B\}} p_{xy}(\alpha(x), \beta(y)),$$

and likewise M_{BC} and M_{CA} . Finally, explicitly calculate

$$M = p_\emptyset M_{AA} M_{AB} M_{BB} M_{BC} M_{CC} M_{CA},$$

the multiplication by p_\emptyset being a scalar product and the others matrix products. Triples (α, β, γ) are in one to one correspondence with assignments $\sigma : V \mapsto [k]$, so the terms of the trace of M ,

$$\text{tr}(M) = \sum_{\alpha} \sum_{\beta} \sum_{\gamma} p_\emptyset M_{AA}(\alpha, \alpha) M_{AB}(\alpha, \beta) M_{BB}(\beta, \beta) M_{BC}(\beta, \gamma) M_{CC}(\gamma, \gamma) M_{CA}(\gamma, \alpha),$$

are in one to one correspondence with the terms of $I(\sigma)$, merely distinguishing whether each monadic score is associated with A , B , or C , and each dyadic score with $A \times A$, $A \times B$, etc. Thus, the partition function is the trace,

$$Z_I = \text{tr}(M).$$

6.2. Complexity. Given that two $n \times n$ matrices over an arbitrary ring can be multiplied using $O(n^\omega)$ ring operations, we have established the following result.

Theorem 10. *Let R be a ring and let G be a graph with n vertices. Let I be any RCSP over R , with constraint graph G and domain $[k]$. Then the algorithm above calculates the partition function Z_I with $O(k^{\omega n/3})$ ring operations.*

For PCSPs, as we discuss in Section 10, under modest conditions each ring operation takes time $O^*(1)$, that is, time polynomial in the input size. In this case, the algorithm will run in time and space $O^*(k^{\omega n/3})$. (Space $O(k^{2n/3})$ suffices if the fast matrix multiplication algorithm uses linear space.)

6.3. A remark. This version of the algorithm is simpler than Williams' because matrix multiplication provides exactly the sum-of-products that RCSP calls for, not the maximum-of-sums needed for the usual 2-CSP. To get the latter, Williams combines the monadic scores with the dyadic ones (this would correspond to incorporating M_{AA} into M_{AB}); "guesses" an optimal score s and how it is partitioned as $s = s_{AB} + s_{BC} + s_{CA}$ among $A \times B$, $B \times C$, and $C \times A$; defines zero-one matrices so that $M_{AB}(\alpha, \beta) = 1$ if the assignments α and β produce the desired value s_{AB} , and 0 otherwise; and multiplies these binary matrices to find (and count) triangles, which correspond to assignments (α, β, γ) yielding $s_{AB} + s_{BC} + s_{CA}$. Williams' original algorithm must iterate over all guesses, of which there are $\Theta(m^3)$ even for 0-1 score functions, while no such guessing is involved in the sum-product version.

7. A REDUCTIVE ALGORITHM

The preponderance of algorithms for CSPs such as Max 2-Sat work by reduction: they reduce an input instance to one or more “smaller” instances in such a way that the solutions to the latter, found recursively, yield a solution to the former. The fastest polynomial-space algorithm for general Max 2-CSP is the reductive algorithm of [SS07], running in time $O^*(k^{19m/100})$ for an instance with m 2-variable constraints and domain $[k]$. While we will not describe that paper’s “Algorithm B” in detail, if an instance I has (a constraint graph with) a vertex of degree 0, 1, or 2, a Type 0, 1, or 2 reduction (respectively) replaces I with an equivalent, smaller instance I' . Otherwise, a Type 3 reduction replaces I with k smaller instances I_0, \dots, I_{k-1} , the largest of whose solutions is the solution to I . If a Type 3 reduction results in a disconnected graph, the components are solved separately and summed.

The algorithm extends to PCSP (and indeed RCSP), the key point being that the Type 0, 1, 2, and 3 CSP reductions have PCSP (and RCSP) analogues. By working with polynomials rather than real numbers we are able to carry substantially more information: for instance, the old CSP reductions correspond to the new PCSP reductions with the polynomials truncated to their leading terms.

7.1. Complexity. We exhibit the extended reductions below, where it will also be evident that each can be performed using $O(k^3)$ ring operations. The extension of [SS07]’s Algorithm B to RCSPs can then be seen to satisfy the following.

Theorem 11. *Let R be a ring and let G be a graph with n vertices and m edges. Let I be any RCSP over R , with constraint graph G and domain $[k]$. Then the extended Algorithm B calculates the partition function Z_I in polynomial space and with $O^*(k^{19m/100})$ ring operations.*

7.2. The reductions. A Type 0 reduction expresses the partition function of an RCSP as a product of partition functions of two smaller instances (or in an important special case, a single such instance). Type 1 and 2 reductions each equate the generating function of an instance to that of an instance with one vertex less. Finally, a Type 3 reduction produces k instances, whose partition functions sum to the partition function of the original instance. In each case, once the reduction is written down, verifying its validity is just a matter of checking a straightforward identity.

A word of intuition, deriving from the earlier CSP reductions, may be helpful (though some readers may prefer to go straight to the equations). A CSP 2-reduction was performed on a vertex v of degree 2, with neighbors u and w . The key observation was that in a maximization problem, the optimal assignment $\sigma(v)$ is a function of the assignments $\sigma(u)$ and $\sigma(w)$. That is, given any assignments $\sigma(u) = i$ and $\sigma(w) = j$, the optimal total $\tilde{s}_{uw}(i, j)$ of the scores of the vertex v , the edges uv and vw , and the edge uw (if present, and otherwise taken to be the 0 score function) is

$$(8) \quad \tilde{s}_{uw}(i, j) = \max_l \{s_{uw}(i, j) + s_{uv}(i, l) + s_v(l) + s_{vw}(l, j)\}.$$

By deleting v from the CSP instance and replacing the original score s_{uw} with the score \tilde{s}_{uw} , we obtain a smaller instance with the same maximum value. A 2-reduction in the RCSP context works the same way, except that the max-of-sums of (8) is replaced by the sum-of-products of (10).

Type 0 Reduction. Suppose I is an RCSP instance whose constraint graph G is disconnected. Let $V = V_1 \cup V_2$ be a nontrivial partition such that $e(V_1, V_2) = 0$. Let I_1 and I_2 be the subinstances obtained by restriction to V_1 and V_2 , except that we define the nullary scores by

$$p_{\emptyset}^{(I_1)} = p_{\emptyset}^{(I)}, \quad p_{\emptyset}^{(I_2)} = 1.$$

A straightforward calculation shows that, for any assignment $\sigma : V \rightarrow [k]$, we have

$$\begin{aligned} I(\sigma) &= p_\emptyset \cdot \prod_{v \in V} p_v(\sigma(v)) \cdot \prod_{xy \in E} p_{xy}(\sigma(x), \sigma(y)) \\ &= \left(p_\emptyset \cdot \prod_{v \in V_1} p_v(\sigma(v)) \cdot \prod_{xy \in E(G[V_1])} p_{xy}(\sigma(x), \sigma(y)) \right) \left(1 \cdot \prod_{v \in V_2} p_v(\sigma(v)) \cdot \prod_{xy \in E(G[V_2])} p_{xy}(\sigma(x), \sigma(y)) \right) \\ &= I_1(\sigma_1) I_2(\sigma_2), \end{aligned}$$

where σ_i denotes the restriction of σ to V_i . It follows easily that

$$Z_I = Z_{I_1} Z_{I_2}.$$

Thus in order to calculate Z_I it suffices to calculate Z_{I_1} and Z_{I_2} .

When one component of G is an isolated vertex v , with $V' = V \setminus v$, one term becomes trivial: $Z_I = (p_\emptyset \cdot \sum_{i \in [k]} p_v(i)) \cdot Z_{I_2}$. So in this case

$$Z_I = Z_{\tilde{I}},$$

where \tilde{I} is the instance obtained from I by deleting v , defining

$$p_\emptyset(\tilde{I}) = p_\emptyset \cdot \sum_{i \in [k]} p_v(i)$$

(a trivial calculation), and leaving all other scores unchanged. This reduces the computation of the partition function of I to that of an instance one vertex smaller.

Type 1 Reduction. Suppose that I is an instance with constraint graph G , and $v \in V$ has degree 1. Let w be the neighbor of v . We shall replace I by an “equivalent” instance \tilde{I} (one with the same partition function) with constraint graph $G \setminus v$.

We define the instance \tilde{I} by giving w vertex scores

$$(9) \quad \tilde{p}_w(i) = p_w(i) \sum_{j=0}^{k-1} (p_{wv}(i, j) \cdot p_v(j)).$$

All other scores remain unchanged except for $p_v(*)$ and $p_{vw}(*, *)$, which are deleted along with v .

To show that $Z_I = Z_{\tilde{I}}$, let $\sigma : V \setminus v \rightarrow [k]$ be any assignment and, for $j \in [k]$, extend σ to $\sigma^j : V \rightarrow [k]$ defined by $\sigma^j(v) = j$ and $\sigma^j|_{V \setminus v} = \sigma$. Using (9), we have

$$\begin{aligned} \tilde{I}(\sigma) &= \tilde{p}_\emptyset \cdot \prod_{x \in V \setminus v} \tilde{p}_x(\sigma(x)) \cdot \prod_{xy \in E(G \setminus v)} \tilde{p}_{xy}(\sigma(x), \sigma(y)) \\ &= \left(p_\emptyset \cdot \prod_{x \in V \setminus \{v, w\}} p_x(\sigma(x)) \cdot \prod_{xy \in E(G \setminus v)} p_{xy}(\sigma(x), \sigma(y)) \right) \cdot \left(p_w(\sigma(w)) \cdot \sum_{j=0}^{k-1} p_{wv}(\sigma(w), j) \cdot p_v(j) \right) \\ &= \sum_{j=0}^{k-1} p_\emptyset \cdot \prod_{v \in V} p_v(\sigma^j(v)) \cdot \prod_{xy \in E} p_{xy}(\sigma^j(x), \sigma^j(y)) \\ &= \sum_{j=0}^{k-1} I(\sigma^j) \end{aligned}$$

and so

$$\begin{aligned}
Z_I &= \sum_{\sigma: V \rightarrow [k]} I(\sigma) \\
&= \sum_{\sigma: V \setminus v \rightarrow [k]} \sum_{j=0}^{k-1} I(\sigma^j) \\
&= \sum_{\sigma: V \setminus v \rightarrow [k]} \tilde{I}(\sigma) \\
&= Z_{\tilde{I}}.
\end{aligned}$$

Type 2 Reduction. Suppose that I is an instance with constraint graph G , and $v \in V$ has degree 2. Let u and w be the neighbors of v in G . We define an instance \tilde{I} with constraint graph \tilde{G} , which will have fewer vertices and edges than G . \tilde{G} is obtained from G by deleting v and adding an edge uw (if the edge is not already present). By analogy with the plain CSP 2-reduction from [SS07] and recapitulated above around (8), we define \tilde{I} by setting, for $i, j \in [k]$,

$$(10) \quad \tilde{p}_{uw}(i, j) = p_{uw}(i, j) \cdot \sum_{l=0}^{k-1} p_{uv}(i, l) p_v(l) p_{vw}(l, j),$$

where we take $p_{uw}(i, j) = 1$ if edge uw was not previously present. All other scores remain unchanged except for $p_v(*)$, $p_{vu}(*, *)$ and $p_{vw}(*, *)$, which are deleted.

To show that $Z_I = Z_{\tilde{I}}$, let $\sigma : V \setminus v \rightarrow [k]$ be any assignment. As before, we write σ^l for the assignment with $\sigma^l|_{V \setminus v} \equiv \sigma$ and $\sigma^l(v) = l$. Then

$$\begin{aligned}
\tilde{I}(\sigma) &= \tilde{p}_{\emptyset} \cdot \prod_{x \in V \setminus v} \tilde{p}_x(\sigma(x)) \cdot \prod_{xy \in E(\tilde{G})} \tilde{p}_{xy}(\sigma(x), \sigma(y)) \\
&= p_{\emptyset} \prod_{x \in V \setminus v} p_x(\sigma(x)) \cdot \prod_{xy \in E(G \setminus v)} p_{xy}(\sigma(x), \sigma(y)) \cdot \sum_{l=0}^{k-1} p_{uv}(\sigma(u), l) p_v(l) p_{vw}(l, \sigma(w)) \\
&= \sum_{l=0}^{k-1} p_{\emptyset} \prod_{x \in V \setminus v} p_x(\sigma^l(x)) \cdot \prod_{xy \in E(G \setminus v)} p_{xy}(\sigma^l(x), \sigma^l(y)) \\
&= \sum_{l=0}^{k-1} I(\sigma^l).
\end{aligned}$$

As with Type 1 reductions, this implies that $Z_I = Z_{\tilde{I}}$.

Type 3 Reduction. Suppose that I is an instance with constraint graph G , and $v \in V$ has degree 3 or more. Let U be the set of neighbors of v in G . We define k instances $\tilde{I}_0, \dots, \tilde{I}_{k-1}$ each with constraint graph $\tilde{G} = G \setminus v$. For $i \in [k]$, the i th instance \tilde{I}_i corresponds to the set of assignments where we take $\sigma(v) = i$.

We define nullary scores for \tilde{I}_i by setting

$$\tilde{p}_{\emptyset}^{(i)} = p_{\emptyset} \cdot p_v(i)$$

and, for each neighbor $u \in U$ of v , and each $j \in [k]$, vertex scores

$$(\tilde{p}^{(i)})_u(j) = p_u(j) \cdot p_{uv}(j, i).$$

All other scores remain unchanged from I except for $p_v(*)$ and $p_{v*}(*, *)$, which are deleted along with v .

For any assignment $\sigma : V \setminus v \rightarrow [k]$, and $i \in [k]$ we write (as usual) $\sigma^i : V \rightarrow [k]$ for the assignment with $\sigma^i|_{V \setminus v} \equiv \sigma$ and $\sigma^i(v) = i$. Then, writing $W = V \setminus (v \cup \Gamma(v))$,

$$\begin{aligned}
\tilde{I}_i(\sigma) &= \tilde{p}_\emptyset \cdot \prod_{x \in V \setminus v} \tilde{p}_x(\sigma(x)) \cdot \prod_{xy \in E(G \setminus v)} \tilde{p}_{xy}(\sigma^i(x), \sigma^i(y)) \\
&= \left(p_\emptyset p_v(i) \right) \cdot \left(\prod_{x \in W} p_x(\sigma(x)) \cdot \prod_{u \in U} [p_u(\sigma(u)) p_{uv}(\sigma(u), i)] \right) \cdot \left(\prod_{xy \in E(G \setminus v)} p_{xy}(\sigma(x), \sigma(y)) \right) \\
&= p_\emptyset \prod_{x \in V} p_x(\sigma^i(x)) \cdot \prod_{xy \in E} p_{xy}(\sigma^i(x), \sigma^i(y)) \\
&= I(\sigma^i).
\end{aligned}$$

Then

$$\begin{aligned}
Z_I &= \sum_{\sigma: V \rightarrow [k]} I(\sigma) \\
&= \sum_{\sigma: V \setminus v \rightarrow [k]} \sum_{i=0}^{k-1} I(\sigma^i) \\
&= \sum_{\sigma: V \setminus v \rightarrow [k]} \sum_{i=0}^{k-1} \tilde{I}_i(\sigma) \\
&= \sum_{i=0}^{k-1} \sum_{\sigma: V \setminus v \rightarrow [k]} \tilde{I}_i(\sigma) \\
&= \sum_{i=0}^{k-1} Z_{\tilde{I}_i}.
\end{aligned}$$

Thus the partition function for I is the sum of the partition functions for the \tilde{I}_i .

We can also write down another reduction, generalizing the Type 0 reduction above, although we will not employ it here.

Cut Reduction. Suppose that I is an instance with constraint graph G , and let $V_0 \subset V$ be a vertex cut in G . Let G_1, \dots, G_r be the components of $G \setminus V_0$. We can calculate Z_I by dividing assignments into classes depending on their restriction to V_0 . Indeed, for each assignment $\sigma_0 : V_0 \rightarrow [k]$, let us define

$$I_0 = p_\emptyset \cdot \prod_{v \in V_0} p_v(\sigma_0(v)) \prod_{xy \in E(G[V_0])} p_{xy}(\sigma_0(x), \sigma_0(y)),$$

and r instances $I_1^{\sigma_0}, \dots, I_r^{\sigma_0}$ where $I_t^{\sigma_0}$ has score polynomials

$$\begin{aligned}
&\tilde{p}_\emptyset = 1 \\
&\forall v \notin V_0 \quad \tilde{p}_v(i) = p_v(i) \cdot \prod_{w \in V_0} p_{vw}(\sigma_0(w), i) \\
&\forall vw \in E[G \setminus V_0] \quad \tilde{p}_{vw}(i, j) = p_{vw}(i, j),
\end{aligned}$$

where we take $p_{vw}(i, j) = 1$ if $vw \notin E$.

Then for any $\sigma : V \rightarrow [k]$, we have

$$I(\sigma) = I_0(\sigma|V_0) \cdot \prod_{i=1}^r I_i^{\sigma|V_0}(\sigma|V_i)$$

and so

$$Z_I = \sum_{\sigma_0: V_0 \rightarrow [k]} I_0(\sigma_0) \prod_{i=1}^r Z_{I_i^{\sigma_0}}.$$

Thus the partition function Z_I is the sum of $|V_0|^k$ products: if $|V_0|$ is small, this provides an effective reduction to smaller cases. Note that this is what we have done in the Type 2 reduction above, where we also use the fact that, for a cutset of size two, the partition functions for the subinstance consisting of the single vertex v can be encoded using constraints on its neighbors. Unfortunately, this is not in general possible for larger cutsets.

8. AN ALGORITHM FOR SPARSE SEMI-RANDOM INSTANCES

An algorithm closely related to the CSP Algorithm B referenced above, and using the same reductions, solves sparse boolean CSP instances in expected linear time [SS06b]. Specifically, for $\lambda = \lambda(n) \geq 0$ and a random graph $G(n, c/n)$ with $c \leq 1 + \lambda n^{-1/3}$ (below the giant-component threshold or in its so-called scaling window), a boolean CSP instance I with constraint graph G is solved in expected time $O(n) \exp(1 + \lambda^3)$.

8.1. Complexity. Because the reductions extend to RCSP, it follows the algorithm extends immediately to boolean RCSP instances on such a constraint graph.

Theorem 12. *Let $\lambda = \lambda(n) \geq 0$. For $c \leq 1 + \lambda n^{-1/3}$, let $G \in \mathcal{G}(n, c/n)$ be a random graph, and let I be any boolean RCSP over any ring R with constraint graph G . Then the algorithm above calculates the partition function Z_I with an expected $O(n) \exp(1 + \lambda^3)$ ring operations.*

9. DYNAMIC PROGRAMMING FOR GRAPHS OF SMALL TREewidth

Roughly, a *tree decomposition* of a graph $G = (V, E)$ consists of a tree T each of whose vertices is associated with a “bag” $B \subseteq V$, such that every vertex of G appears in some bag, every edge of G has both endpoints in some bag, and the set of bags containing any given vertex v of G forms a subtree of T . The “width” of a tree decomposition is one less than the cardinality of the largest bag, and the treewidth of G is the minimum width of any tree decomposition. For our purposes we will assume that a tree decomposition is given, though for graphs of constant-bounded treewidth, a minimum-treewidth decomposition can be found in linear time [RS95, Ree92, Bod96].

Efficient algorithms for various sorts of constraint satisfaction and related problems on graphs of small treewidth have been studied since at least the mid-1980s, with systematic approaches dating back at least to [DP87, DP89, AP89]. A special issue of Discrete Applied Mathematics was devoted to this and related topics in 1994 [AHE94], and the field remains an extremely active area of research.

A recent series of results on treewidth is Monien and Preis’s proof that every cubic graph with m edges has bisection width at most $(1/6 + o(1))m$ [MP01], Fomin and Høie’s use of this to show that the pathwidth of a cubic graph is bounded by $(1/6 + o(1))m$ [FH06], and the use of this to prove that any graph has treewidth at most $(13/75 + o(1))m$ [SS07, KMRR09].

Focusing on two problems of particular interest to us, as a key part of an approximation result in Jansen, Karpinski, Lingas and Seidel’s [JKLS05], it was shown that given a graph G and a tree decomposition T of width $b - 1$, an exact maximum (or minimum) bisection of G can be found in time *and space* $O^*(2^b)$, using dynamic programming on the tree decomposition. To find a maximum

clique in time $O^*(2^b)$ is even easier, and may be done in polynomial space, using the well-known fact that any clique of G is necessarily contained in a single bag of T .

Here, we generalize these results for bisection and clique to the entire class RCSP. While the extension is reasonably straightforward, the details are confusing enough that it is valuable to have the general result available in this neatly packaged form, and not have to apply the approach (or wonder if it can be applied) to a particular problem.

9.1. Complexity.

Theorem 13. *Let G be a graph with n vertices, m edges, and a tree decomposition T of width $b-1$. Let I be any RCSP instance over G . Given I and T , the partition function Z_I can be calculated in space $O(k^{b-1})$ and with $O^*(k^b)$ ring operations.*

We sketch the proof before detailing it. We express an assignment score as a product of “bag scores”, f_B for a bag B . In computing the partition function, we reduce on any leaf bag B_1 of T and its parent B_2 , in a way closely related to the Type 1 reductions of Section 7: variables appearing *only* in B_1 can be “integrated out” to yield a function f'_{B_1} on a smaller set of variables. By the definition of a tree decomposition, that smaller set is a subset of the variables of B_2 , so we can absorb f'_{B_1} into f_{B_2} by defining $f'_{B_2} = f_{B_2} \cdot f'_{B_1}$. Deleting B_1 from T , and forgetting f_{B_1} , completes the reduction. We now present this more formally.

Proof. We first show how to express the score of an assignment as a product of “bag scores”. Let \mathcal{B} be the collection of all bags. While a vertex is typically found in several bags, associate each vertex $v \in V(G)$ with just one bag that contains it; similarly, associate each edge $e \in E(G)$ with some bag containing both its endpoints. For a bag B , let $V(B)$ be the set of associated vertices and $E(B)$ the set of associated edges, so that $\uplus_{B \in \mathcal{B}} V(B) = V(G)$ and $\uplus_{B \in \mathcal{B}} E(B) = E(G)$. Define each bag score as the product of its associated vertex and edge scores:

$$(11) \quad f_B(\sigma|_B) = \prod_{v \in V(B)} p_v(\sigma(v)) \cdot \prod_{xy \in E(B)} p_{xy}(\sigma(x), \sigma(y)).$$

(Note that for an edge $xy \in E(B)$, the endpoints x and y need not be in $V(B)$, but they must be in B itself, and thus $f_B(\sigma|_B)$ is well defined, depending only on assignments of vertices in B .) It is clear that for the original PCSP instance I ,

$$(12) \quad \begin{aligned} Z_I &:= \sum_{\sigma: V \rightarrow [k]} p_\emptyset \cdot \prod_{v \in V} p_v(\sigma(v)) \cdot \prod_{xy \in E} p_{xy}(\sigma(x), \sigma(y)) \\ &= \sum_{\sigma: V \rightarrow [k]} p_\emptyset \cdot \prod_{B \in \mathcal{B}} f_B(\sigma|_B). \end{aligned}$$

We now show a reduction that will preserve the form of (12), but with one bag fewer. After reduction, the bag scores themselves will no longer depend simply on vertex and edge scores as in (11), which is why we work with (12) instead.

Let B_1 be a leaf bag and B_2 its parent. Let $V' = B_1 \setminus B_2$. By the nature of a tree decomposition, V' is the set of vertices found exclusively in B_1 . Let $V'' = V \setminus V'$. Correspondingly, we will use σ' and σ'' for assignments from these sets to $[k]$, and (σ', σ'') for a complete assignment from V to $[k]$. In the following rewriting of (12), lines (14) and (15) are implicit definitions of the functions f'_{B_1}

and f'_{B_2} .

$$\begin{aligned}
(13) \quad Z_I &= \sum_{\sigma: V \rightarrow [k]} p_\emptyset \cdot \left(\prod_{B \in \mathcal{B} \setminus \{B_1, B_2\}} f_B(\sigma|_B) \right) \cdot f_{B_2}(\sigma|_{B_2}) \cdot f_{B_1}(\sigma|_{B_1}) \\
&= \sum_{\sigma'': V'' \rightarrow [k]} p_\emptyset \cdot \left(\prod_{B \in \mathcal{B} \setminus \{B_1, B_2\}} f_B(\sigma''|_B) \right) \cdot f_{B_2}(\sigma''|_{B_2}) \cdot \sum_{\sigma': V' \rightarrow [k]} f_{B_1}((\sigma'', \sigma')|_{B_1}) \\
(14) \quad &= \sum_{\sigma'': V'' \rightarrow [k]} p_\emptyset \cdot \left(\prod_{B \in \mathcal{B} \setminus \{B_1, B_2\}} f_B(\sigma''|_B) \right) \cdot f_{B_2}(\sigma''|_{B_2}) \cdot f'_{B_1}(\sigma''|_{B_1 \cap B_2}) \\
(15) \quad &= \sum_{\sigma'': V'' \rightarrow [k]} p_\emptyset \cdot \left(\prod_{B \in \mathcal{B} \setminus \{B_1, B_2\}} f_B(\sigma''|_B) \right) \cdot f'_{B_2}(\sigma''|_{B_2}).
\end{aligned}$$

This reduces the number of bags by one, as desired. Note that the definition implicit in (14), that for any $\sigma'': V'' \rightarrow [k]$,

$$f'_{B_1}((\sigma'')|_{B_1 \cap B_2}) := \sum_{\sigma': V' \rightarrow [k]} f_{B_1}((\sigma'', \sigma')|_{B_1}),$$

is well defined because the right-hand side depends only on values of (σ'', σ') on the set $B_1 \cap B_2$. Likewise, the definition implicit in (15), that for any $\sigma'': V'' \rightarrow [k]$,

$$f'_{B_2}(\sigma''|_{B_2}) := f_{B_2}(\sigma''|_{B_2}) \cdot f'_{B_1}((\sigma'')|_{B_1 \cap B_2}),$$

is well defined because the right-hand side depends only on values of (σ'', σ') on the set B_2 .

In going from (13) to (15) we reduced the decomposition tree size by 1. To do so we represented the functions f'_{B_1} and f'_{B_2} explicitly, writing down respectively $k^{|V'|}$ and $k^{|V(B_2)|}$ score values. The reduction does not change the size of any bag, so if initially each bag has size at most $b-1$, each reduction takes space $O(k^{b-1})$ and requires $O^*(k^{b-1})$ ring operations (each bag evaluation may require evaluating polynomially many vertex and edge scores).

The algorithm really does require extra space $O(k^{b-1})$ to represent the updated score function f'_{B_2} , because f'_{B_2} incorporates f_{B_1} “integrated out” over $\sigma': V' \rightarrow [k]$ and can no longer be viewed as coming from a bag of vertex and edge scores. We conclude that altogether the algorithm takes space $O^*(k^{b-1})$ and uses $O^*(k^{b-1})$ ring operations. \square

10. CONTROLLING THE RING-OPERATION COMPLEXITY

In this section, we restrict to Polynomial CSPs, and consider the complexity of performing the ring operations in our algorithms. In general, PCSP instances can be intractable. For example, consider a boolean ($[k] = \{0, 1\}$) PCSP instance I where all dyadic scores are 1, and the monadic scores are $p_v(0) = 1$, $p_v(1) = z^{r_v}$, with mutually incommensurable real values r_v . Then each of the 2^n variable assignments corresponds to a different monomial in the partition function $Z(I)$. Since the output is of size $\Omega(2^n)$, there is no hope of a time- or space-efficient algorithm.

However, some modest restrictions suffice to control the size of the elements of the polynomial ring with which we are working (including the output itself), and thus the time- and space-efficiency of the various algorithms.

10.1. Polynomially bounded instances. The simplest well-behaved case is when an instance is over a single variable, and all score polynomials are of low degree, with small coefficients. This is easily formalized and generalized.

Definition 14 (Polynomially bounded). *A family \mathcal{F} of instances of PCSP is polynomially bounded if for an instance $I \in \mathcal{F}$ having n variables and m clauses, with respect to $m+n$:*

- the number of formal variables is $O(1)$,
- all powers are integers with absolute value $O^*(1)$, and
- coefficients are integers of length $O^*(1)$, or sums and products of coefficients are deemed to be elementary operations.

For convenience we will speak of “a polynomially bounded instance”.

Lemma 15. *In any of the four algorithms described, over a fixed domain $[k]$, solving a polynomially bounded instance I with n variables and m clauses, each ring operation can be performed in time and space $O^*(1)$.*

Proof. Say that there are c variables, each with maximum degree Δ . As a sum of products, the partition function itself has maximum degree (in each variable) at most $(m + n + 1)\Delta$. This is also true of any polynomial that arises in the course of the computation. This can be verified directly by looking at the details of each algorithm. Alternatively, note that the algorithms compute “obliviously” on the polynomials (the operations performed depend on the constraint graph but not on the polynomials), the algorithms never subtract, and thus if on any input an algorithm ever produced terms of higher degree than the degree bound on the partition function, it would certainly do so when each input score polynomial is $1 + z + \dots + z^\Delta$ (or the equivalent multivariate polynomial), and (in the absence of any negative coefficients) such terms would survive to appear in the partition function, a contradiction.

Since each intermediate polynomial has maximum degree at most $(m + n + 1)\Delta$, it has at most $((m + n + 1)\Delta)^c = O^*(1)$ terms, and two such polynomials can be multiplied with $O^*(1)$ integer operations.

If we wish to further break down the time as a function of the lengths of the integer coefficients, each arises as a sum of at most k^{m+n+1} products of at most $m + n + 1$ input coefficients, and therefore has length at most $(m + n + 1) \log k \cdot (m + n + 1) = O^*(1)$ times that of the longest input coefficient. Thus, adding and multiplying coefficients can also be done in time $O^*(1)$. \square

10.2. Pruned instances. An alternative for a PCSP in a single variable z is to allow arbitrary (real) powers in the score polynomials, but to avoid the blowup in our motivating example by demanding not the full partition function but just its leading term (the one with the highest power of z).

Definition 16 (z -prunable). *An instance of PCSP with n variables and m clauses (more precisely a family of instances) is z -prunable if, with respect to $m + n$:*

- the number of formal variables is $O(1)$,
- z may have arbitrary real powers,
- the remaining variables w_1, w_2, \dots have integer powers with absolute value no more than $O^*(1)$, and
- coefficients are nonnegative integers of length $O^*(1)$, or coefficients are nonnegative reals and sums and products of coefficients are deemed to be elementary operations,

where the $O(\cdot)$ notation is with respect to $m + n$.

Note that a PCSP arising as the generating function of a CSP (see Section 5.1) is always z -prunable: there is just a single formal variable z , and all coefficients are 1.

Pruned polynomial. Given a polynomial p in one variable z , we define the *pruned polynomial* $(p)_z$ to be the polynomial obtained by removing all but the leading term. If p is a polynomial in variables z, w_1, w_2, \dots , we obtain $(p)_z$ by throwing away all terms T such that there is a term of the form cTz^i , where $c > 0$ is real, and $i > 0$. For instance,

$$(2z^2 + 3z + 700 + zw_1 + z^2w_1 + zw_2 + z^{10}w_1w_2)_z = 2z^2 + z^2w_1 + zw_2 + z^{10}w_1w_2.$$

Pruning a partition function often preserves all the information of interest. Consider for example a PCSP I arising from a simple maximization CSP, say weighted Max Cut. In the partition function Z_I , the leading term's power is the weight of a largest cut, and its coefficient the number of such cuts. (For this informal discussion, we will simply multiply-count cuts that are symmetric to one another.) With real edge weights, Z_I can have exponentially many terms, but the weight of a largest cut, and the number of such cuts, is preserved in the pruned partition function $(Z_I)_z$, which consists *solely* of that leading term.

Somewhat more generally, consider a CSP where we wish to maximize one (real-valued) parameter for one or more values of several other (integer-valued) parameters. For example, the maximum bisection of an edge-weighted graph means maximizing the weight of cut edges, when the number of vertices in one partition is held to $n/2$. Terms in the partition function Z_I of the corresponding bivariate PCSP have real powers of the edge-weighting variable z , and integral powers of the variable w that counts vertices in partition 1. Again (by definition) pruning to $(Z_I)_z$ preserves the maximum power of z and the corresponding coefficient for any fixed power of w , so from $(Z_I)_z$ we may read off the weight of a maximum bisection and the number of such bisections, and indeed the weight and cardinality of maximum cuts of any specified sizes (say with exactly $n/3$ vertices in one partition).

Pruning in the algorithms. The key point about pruning is that for polynomials p and q (in any set of variables) if coefficients are nonnegative then $(p + q)_z = (p_z + q_z)_z$ and $(pq)_z = ((p)_z(q)_z)_z$.

With either of the reductive algorithms described, the pruned partition function can be obtained efficiently, first pruning the input instance, then pruning as we go. (Correctness of this can be seen inductively, working backwards from the last reduction.) It is a trivial observation that for each reduction R (of Type 0–3), and any instance I , $(R(I))_z = (R(I_z))_z$. In particular, if we perform a full sequence of reductions and z -prune at every stage, we will end up with $(Z_I)_z$.

Pruning can also be done for tree decomposition-based dynamic-programming algorithms. Referring to the proof of Theorem 13, we simply note that the pruned partition function can be obtained from pruned versions of f_{B_1} , f'_{B_1} , f_{B_2} , and f'_{B_2} .

This establishes the following lemma.

Lemma 17. *Given a z -prunable PCSP instance I with n variables and m clauses over a fixed domain $[k]$, the pruned partition function can be calculated by the reductive or dynamic programming algorithms (see Theorems 11, 12, and 13) with each ring operation taking time and space $O^*(1)$.*

It is not likely that the PCSP extension of Williams' algorithm can accommodate real-valued powers, even with pruning. One problem is the use of fast matrix multiplication as a “black box”: it is unlikely that we can prune within the multiplication algorithm. Another piece of evidence against being able to accommodate real-valued PCSP powers is that in the CSP setting, Williams was unable to accommodate real-valued scores (the analogue of the PCSP's powers), and a solution for PCSP would imply one for CSP. Williams noted that one could work around real-valued scores by approximation methods, for example multiplying all scores by a large constant and then replacing each with the nearest integer, and this can also be done in the PCSP setting.

11. SUMMARY OF PCSP ALGORITHMIC RESULTS

Combining the results in Sections 6–9, phrased in terms of the number of ring operations, with the bounds in Section 10 on the complexity of these operations for PCSPs, gives the following general conclusions. Some particular results are presented in Section 13.

Theorem 18. *The PCSP extension of Williams split-and-list algorithm [Wil04] solves any polynomially bounded PCSP instance with n variables over domain $[k]$ in time and space $O^*(k^{\omega n/3})$, where ω is the matrix-multiplication exponent for the ring of polynomials over reals.*

Parametrizing by vertices, this exponential-space algorithm is the only known algorithm faster than $O^*(k^n)$ for PCSP, just as it is for CSP. In particular, it is the only algorithm efficient for dense PCSP instances.

Theorem 19. *The PCSP extension of of Scott and Sorkin’s reductive algorithm [SS07] solves any polynomially bounded PCSP instance (or finds the pruned partition function of any prunable PCSP instance) with n variables, m clauses, and over domain $[k]$, in time $O^*(k^{19m/100})$ and space $O^*(1)$.*

This algorithm is suitable for sparse instances (if $m > (100/19)n$ then the naive $O(k^n)$ algorithm is better). It is not as fast as dynamic programming (see Theorem 21), but unlike that and Williams’ algorithm it runs in polynomial space, and is potentially practical.

Theorem 20. *When $G = G(n, c/n)$ is an Erdős-Rényi random graph with $c \leq 1 + \lambda n^{-1/3}$, an extension of Scott and Sorkin’s expected-linear-time CSP algorithm [SS06b] solves any boolean PCSP instance with constraint graph G in expected time $O^*(1) \exp(1 + \lambda^3)$ and in space $O^*(1)$.*

This polynomial-space, polynomial-expected-time algorithm is of course the most efficient for sparse semi-random instances (below the giant-component threshold or in its scaling window).

Theorem 21. *Tree decomposition-based dynamic programming can solve any polynomially bounded PCSP instance (or find the pruned partition function of any prunable PCSP instance) with n variables, m clauses, over domain $[k]$, and having treewidth $b - 1$, in time and space $O^*(k^b)$ or (since $b - 1 \leq (13/75 + o(1))m$, from [SS07, KMRR09]) time and space $O^*(k^{(13/75 + o(1))m})$.*

This exponential-space algorithm is suitable for instances of small treewidth, including sparse instances with $m < (75/13)n$.

12. CONSTRUCTING AND SAMPLING SOLUTIONS

Wherever we can compute a CSP’s maximum value, we can also produce a corresponding assignment; and wherever we can count assignments producing a given value, we can also do exact random sampling from these assignments. The method is standard, and we illustrate with sampling. We construct our assignment one variable at a time, starting from the empty assignment. Given a partial assignment $\sigma_0 : V_0 \rightarrow [k]$, and a vertex $v \notin I_0$, we calculate the partition functions

$$(16) \quad Z_{I; \sigma_0(i)} = \sum_{\sigma : \sigma|_{V_0} = \sigma_0, \sigma(v)=i} I(\sigma),$$

and use these to determine the conditional distribution of $\sigma(v)$ given that $\sigma|_{I_0} = \sigma_0$.

This enables us to sample from a variety of distributions. For instance, we get the following result.

Theorem 22. *Let G be a graph with m edges. Then in time $O^*(2^{19m/100})$ and space $O^*(1)$ we can sample uniformly at random from the following distributions:*

- *maximum cuts, maximum bisections, minimum bisections*
- *maximum independent sets, cliques of maximal size*
- *independent sets of any fixed size, cliques of any fixed size.*

In the same time, we can sample from the equilibrium (Gibbs) distribution of the Ising model with any fixed interaction strength and external magnetic field.

A similar result holds for edge- and/or vertex-weighted graphs, except that we sample at random only from optimal assignments. Many other problems can be expressed in this framework. For example, we can count and sample proper k -colorings in time $O^*(k^{19m/100})$ and space $O^*(1)$, or in time and space $O^*(k^{(13/75 + o(1))m})$.

If we wish to exhibit just one optimal assignment deterministically, we can do so by a small modification of the sampling approach above: at each step, assign the next vertex v the smallest i such that $\sigma_0(i)$ can be extended to an optimal assignment (this can be read off from $Z_{I;\sigma(i)}$ given by (16)).

13. APPLICATIONS AND CONCLUSIONS

The PCSP / RCSP framework encompasses many problems that are not in CSP, yet can be solved by extensions of any general CSP algorithm we know of, with essentially equal efficiency. For several problems, this yields the best algorithms known.

13.1. Graph bisection. Jansen, Karpinski, Lingas and Seidel [JKLS05] showed how to use dynamic programming based on tree decomposition to find a maximum or minimum bisection of a graph. Combined with the bounds from [SS07, KMRR09] on the treewidth of a graph with m edges, this gave an algorithm with a time and space bound of $O^*(2^{(13/75+o(1))m})$.

Until now, no polynomial-space algorithm was known for this problem other than the naive enumeration running in time $O^*(2^n)$. Once bisection is phrased as a PCSP, our results (Theorem 19) immediately imply that in time $O^*(2^{19m/100})$ and polynomial space, we can solve maximum and minimum bisection, count maximum and minimum bisections, and indeed count bisections of all sizes. We can do this also in weighted bisection models, as long as they are polynomially bounded or z -prunable.

Our results (Theorem 18) also mean that dense instances of bisection can be solved by Williams' algorithm, in time and space $O^*(2^{\omega n/3})$, again yielding the full generating function of cuts.

13.2. Cliques and Independent Sets. Maximum Independent Set is in Max 2-CSP and has been studied extensively, but, surprisingly, PCSP can offer an advantage here. For simple MIS, for any graph with average vertex degree d , all four of our general-purpose PCSP algorithms are dominated either by [BEP08] (for average degree ≤ 3 , time $O^*(2^{0.1331n})$) or by [FGK06] (time $O^*(2^{0.288n})$, otherwise). However, for weighted MIS, counting MIS, or weighted counting MIS, the fastest algorithm we are aware of is the vertex-parametrized $O^*(2^{0.3290n})$ algorithm of Fürer and Kasiviswanathan [FK05], based on an improved analysis of an algorithm of Dahllöf, Jonsson, and Wahlström [DJW05]. For any of these problems the $O^*(2^{19m/100})$ reductive algorithm of Theorem 19 is the fastest polynomial-space algorithm known if the average degree is below about 3.46, and the $O^*(2^{(13/75+o(1))m})$ exponential-space dynamic-programming algorithm of Theorem 21 the fastest algorithm known if the average degree is below about 3.79.

We noted earlier that Maximum Clique is not in the class 2-CSP, because constraints are given by non-edges, not edges, but that it is in the class PCSP. This is nice, but does not currently offer any algorithmic advantage. If G has average degree $d = o(n)$, the following simple algorithm has running time $O^*(2^{o(n)})$, dominating all other algorithms under consideration here. For any k of our choosing, any clique either is contained in the $\leq n/k$ -order subgraph induced by vertices of degree $\geq kd$, which is checkable in time $O^*(2^{n/k})$, or contains a vertex of degree $\leq kd$, which is checkable in time $O^*(2^{kd})$. Taking $k = \sqrt{n/d}$ gives running time $O^*(2^{\sqrt{nd}})$. If d is not $o(n)$ then certainly $d = \omega(1)$. This rules out edge-parametrized algorithms, leaving us only Williams' algorithm, which is less efficient than applying a vertex-parametrized MIS algorithm such as that of [FK05] to the complement graph \overline{G} .

13.3. Ising partition function. The Ising partition function, a canonical object in statistical physics, is a PCSP, and can be solved by any of the algorithms presented here. We know of no other general algorithm to find the partition function other than naive enumeration over all 2^n configurations.

13.4. Potts and other q -state models. The Potts (q -coloring) partition function can trivially be evaluated in time $O^*(q^n)$, but it is by no means obvious that one can remove the dependence on q . A breakthrough result of Björklund, Husfeldt, Kaski, and Koivisto [BHKK08] (see also [BH06b, Koi06a, BHK, BHKK07, BH06a, BH08]) shows how to evaluate the Potts partition function (and therefore also the Tutte polynomial) using the Inclusion–Exclusion method, in time $O^*(2^n)$ in exponential space, or time $O^*(3^n)$ in polynomial space.

Our PCSP model includes the Potts model and more general models where the “energy” is not simply a function of equality or inequality of neighboring colors. Our algorithms unfortunately do have exponential dependence on q , but with the greater generality of PCSP this may be unavoidable: as noted in [BHKK08], Traxler [Tra08] shows that if the q -state Potts model is generalized even to 2-CSPs on variables with q states, exponential dependence on q cannot be avoided, at least in the vertex-parametrized case, assuming the Exponential Time Hypothesis.

13.5. Conclusions. As we have shown, the PCSP formulation is clean and offers quantifiable advantages for a wide range of problems. However, the greatest benefit is its broad applicability. For example, exponential time algorithms have not previously been considered for problems such as judicious partitioning, and for virtually any form of this problem such algorithms now follow instantly from its membership in PCSP. Moreover, when solving optimization problems coded into PCSP form, we automatically get solutions to the corresponding counting problems.

Finally, we note that algebraic formulations appear to be playing a newly important role in the development of exponential-time algorithms for CSPs and related problems. Our work along these lines began with [SS06a], and we have already discussed Koivisto’s sum-of-products variation on Williams’ split-and-list algorithm [Koi06b]. The methods leading to the $O^*(2^n)$ Tutte polynomial computation [BHKK08] work in a fairly general algebraic setting, and a new algorithm of Williams [Wil08] that finds a k -path in an order- n graph in time $O^*(2^k)$ works in an algebra carefully chosen for properties including making walks other than simple paths cancel themselves out. PCSPs should be considered one piece of this new algebraic tool kit.

REFERENCES

- [AHE94] S. Arnborg, S.T. Hedetniemi, and A. Proskurowski (Eds.), *Discrete Applied Mathematics* **54** (1994), no. 2-3, special issue.
- [AP89] S. Arnborg and A. Proskurowski, *Linear time algorithms for NP-hard problems restricted to partial k -trees*, *Discrete Applied Mathematics* **23** (1989), no. 1, 11–24.
- [BEP08] Nicolas Bourgeois, Bruno Escoffier, and Vangelis Th. Paschos, *An $O^*(1.0977^n)$ exact algorithm for max independent set in sparse graphs*, IWPEC (Martin Grohe and Rolf Niedermeier, eds.), *Lecture Notes in Computer Science*, vol. 5018, Springer, 2008, pp. 55–65.
- [BG05] Andrei Bulatov and Martin Grohe, *The complexity of partition functions*, *Theoret. Comput. Sci.* **348** (2005), no. 2-3, 148–186.
- [BH06a] Andreas Björklund and Thore Husfeldt, *Exact algorithms for exact satisfiability and number of perfect matchings*, *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP 2006)*, *Lecture Notes in Computer Science*, vol. 4051, Springer, 2006, pp. 548–559.
- [BH06b] ———, *Inclusion-exclusion algorithms for counting set partitions*, *Proceedings of the 47th IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, 2006, pp. 575–582.
- [BH08] Andreas Björklund and Thore Husfeldt, *Exact algorithms for exact satisfiability and number of perfect matchings*, *Algorithmica* **52** (2008), no. 2, 226–249.
- [BHK] Andreas Björklund, Thore Husfeldt, and Mikko Koivisto, *Set partitioning via inclusion-exclusion*, *SIAM Journal on Computing*, special issue dedicated to selected papers from FOCS 2006, to appear.
- [BHKK07] Andreas Björklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto, *Fourier meets Möbius: fast subset convolution*, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC 2007)*, ACM, New York, 2007, pp. 67–74.
- [BHKK08] Andreas Björklund, Thore Husfeldt, Petteri Kaski, and Mikko Koivisto, *Computing the Tutte polynomial in vertex-exponential time*, *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science (FOCS 2008)*, 2008.

- [Bod96] Hans L. Bodlaender, *A linear-time algorithm for finding tree-decompositions of small treewidth*, SIAM J. Comput. **25** (1996), no. 6, 1305–1317.
- [BS99] Béla Bollobás and Alexander D. Scott, *Exact bounds for judicious partitions of graphs*, Combinatorica **19** (1999), no. 4, 473–486.
- [CW90] Don Coppersmith and Shmuel Winograd, *Matrix multiplication via arithmetic progressions*, J. Symbolic Comput. **9** (1990), no. 3, 251–280.
- [DGP06] Martin E. Dyer, Leslie Ann Goldberg, and Mike Paterson, *On counting homomorphisms to directed acyclic graphs*, ICALP (1), 2006, pp. 38–49.
- [DJW05] Vilhelm Dahllöf, Peter Jonsson, and Magnus Wahlström, *Counting models for 2SAT and 3SAT formulae*, Theoret. Comput. Sci. **332** (2005), no. 1-3, 265–291.
- [DP87] R. Dechter and J. Pearl, *Network-based heuristics for constraint-satisfaction problems*, Artif. Intell. **34** (1987), no. 1, 1–38.
- [DP89] Rina Dechter and Judea Pearl, *Tree clustering for constraint networks (research note)*, Artif. Intell. **38** (1989), no. 3, 353–366.
- [FGK06] Fedor V. Fomin, Fabrizio Grandoni, and Dieter Kratsch, *Measure and conquer: a simple $O(2^{0.288n})$ independent set algorithm*, Proceedings of the 17th annual ACM-SIAM symposium on Discrete Algorithms (SODA 2006) (New York, NY, USA), ACM, 2006, pp. 18–25.
- [FH06] Fedor V. Fomin and Kjartan Høie, *Pathwidth of cubic graphs and exact algorithms*, Inform. Process. Lett. **97** (2006), no. 5, 191–196.
- [FK05] Martin Fürer and Shiva Prasad Kasiviswanathan, *Algorithms for counting 2-SAT solutions and colorings with applications*, Tech. Report 33, Electronic Colloquium on Computational Complexity, 2005.
- [HN04] Pavol Hell and Jaroslav Nešetřil, *Graphs and homomorphisms*, Oxford Lecture Series in Mathematics and its Applications, vol. 28, Oxford University Press, Oxford, 2004.
- [JKLS05] Klaus Jansen, Marek Karpinski, Andrzej Lingas, and Eike Seidel, *Polynomial time approximation schemes for max-bisection on planar and geometric graphs*, SIAM J. Comput. **35** (2005), no. 1, 110–119 (electronic).
- [KMRR09] Joachim Kneis, Daniel Mölle, Stefan Richter, and Peter Rossmanith, *A bound on the pathwidth of sparse graphs with applications to exact algorithms*, SIAM J. Discrete Math., vol. 23, SIAM, 2009, pp. 407–427.
- [Koi06a] Mikko Koivisto, *An $O^*(2^n)$ algorithm for graph colouring and other partitioning problems via inclusion-exclusion*, Proceedings of the 47th IEEE Symposium on Foundations of Computer Science (FOCS 2006), 2006, pp. 583–590.
- [Koi06b] Mikko Koivisto, *Optimal 2-constraint satisfaction via sum-product algorithms*, Inform. Process. Lett. **98** (2006), no. 1, 24–28.
- [MP01] Burkhard Monien and Robert Preis, *Upper bounds on the bisection width of 3- and 4-regular graphs*, Mathematical foundations of computer science, 2001 (Mariánské Lázně), Lecture Notes in Comput. Sci., vol. 2136, Springer, Berlin, 2001, pp. 524–536.
- [Ree92] Bruce A. Reed, *Finding approximate separators and computing tree width quickly*, STOC, ACM, 1992, pp. 221–228.
- [RS95] N. Robertson and P. D. Seymour, *Graph minors XIII: The disjoint path problem*, J. Combinatorial Theory (Series B) **63** (1995), 65–110.
- [Sco05] Alexander D. Scott, *Judicious partitions and related problems*, Surveys in combinatorics 2005, London Math. Soc. Lecture Note Ser., vol. 327, Cambridge Univ. Press, Cambridge, 2005, pp. 95–117.
- [SS06a] Alexander D. Scott and Gregory B. Sorkin, *Generalized constraint satisfaction problems*, Tech. Report cs:DM/0604079v1, arxiv.org, April 2006, See <http://arxiv.org/abs/cs.DM/0604079>.
- [SS06b] ———, *Solving sparse random instances of Max Cut and Max 2-CSP in linear expected time*, Comb. Probab. Comput. **15** (2006), no. 1-2, 281–315.
- [SS07] Alexander D. Scott and Gregory B. Sorkin, *Linear-programming design and analysis of fast algorithms for Max 2-CSP*, Discrete Optim. **4** (2007), no. 3-4, 260–287.
- [Tra08] Patrick Traxler, *The time complexity of constraint satisfaction*, Proceedings of the 3rd International Workshop on Parameterized and Exact Computation (IWPEC 2008), Lecture Notes in Computer Science, vol. 5018, Springer, 2008, pp. 190–201.
- [Wil04] Ryan Williams, *A new algorithm for optimal constraint satisfaction and its implications*, Proc. 31st International Colloquium on Automata, Languages and Programming (ICALP), 2004.
- [Wil08] ———, *Finding paths of length k in $O^*(2^k)$ time*, Tech. Report cs.DS/0807.3026v3, arxiv.org, November 2008, See <http://arxiv.org/abs/cs.DM/0604080>.

(Alexander D. Scott) MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, 24-29 ST GILES', OXFORD, OX1 3LB, UK

E-mail address: `scott@maths.ox.ac.uk`

(Gregory B. Sorkin) DEPARTMENT OF MATHEMATICAL SCIENCES, IBM T.J. WATSON RESEARCH CENTER, YORKTOWN HEIGHTS NY 10598, USA

E-mail address: `sorkin@watson.ibm.com`